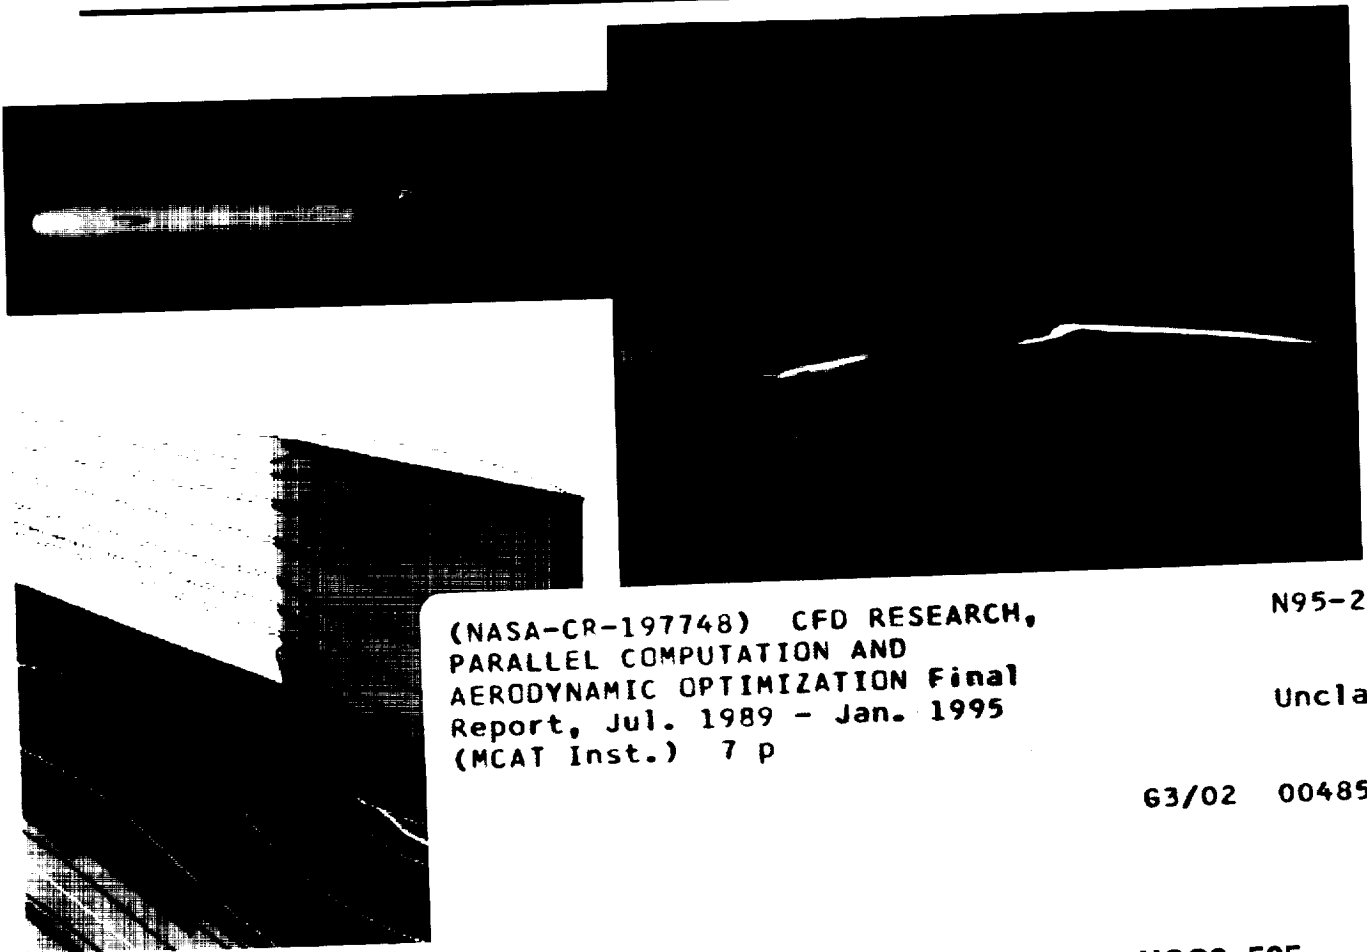


MCAT Institute
Final Report
95-18

CFD Research, Parallel Computation and Aerodynamic Optimization

James S. Ryan



(NASA-CR-197748) CFD RESEARCH,
PARALLEL COMPUTATION AND
AERODYNAMIC OPTIMIZATION Final
Report, Jul. 1989 - Jan. 1995
(MCAT Inst.) 7 p

N95-26649

Unclas

63/02 0048503

January 1995

MCAT Institute
3933 Blue Gum Drive
San Jose, CA 95127

NCC2-505
NCC2-796

FINAL REPORT

CFD Research, Parallel Computation and Aerodynamic Optimization

James S. Ryan
MCAT Institute
NCC 2-505 and 2-796
July 1989 to January 1995

Introduction

This report covers over five years of research in CFD and its applications. During that time, I completed the publication of work on hypersonic flows, modeled incompressible flow about submarine hulls, and computed the flow about supersonic transport aircraft on parallel computers. Finally, using CFD as an established tool, I began to explore aerodynamic optimization on parallel architectures.

The objective of this work has always been to provide better tools to vehicle designers. Submarine design requires accurate force and moment calculations in flows with thick boundary layers and large separated vortices. Low noise production is critical, so flow into the propulsor region must be predicted accurately.

The High Speed Civil Transport (HSCT) has been the subject of the most recent work. This vehicle is to be a passenger aircraft, with the promise of cutting overseas flight times by more than half. A successful design must far surpass the performance of the only existing comparable plane, the Concorde. Fuel economy, other operational costs, environmental impact, and range must all be improved substantially. The aircraft must be able to fly many more routes than the Concorde, and can only do so if noise production can be limited. For all of these reasons, improved design tools are required, and these tools must eventually integrate optimization, external aerodynamics, propulsion, structures, heat transfer, controls, and perhaps other disciplines. If this project contributes to improved design tools for US industry, and thus to our economic competitiveness, it is successful.

The work was done under cooperative agreements NCC 2-505 and NCC 2-796. I was a co-principal investigator under the first agreement from July 1989 to May 1993. I was principal investigator under the second agreement from May 1993 until January 1995. The work completed during these periods will be highlighted in roughly chronological order.

Incompressible Flow for Submarine Hydrodynamics

The goal of this project was to accurately and quickly compute flows about submarine hulls. The results would be used in the DARPA SUBOFF project to evaluate the state of the art in hydrodynamic codes, and to select codes for future use in submarine design.

The initial work was done with the INS-3D code. Accomplishments included moving the code to Cray-2 and Cray Y-MP architectures, addition of the Baldwin-Lomax turbulence model to the code, and implementation of a zonal grid scheme. Solutions were converged on a 10-zone grid, representing the hull, fairwater (sail), and four tail appendages. A 13-zone solution with 1.9 million points and improved boundary layer resolution was in progress when emphasis was shifted to the INS3D-LU code.

In a meeting at Ames in December 1989⁷, representatives of the SUBOFF project emphasized the need for advanced turbulence models and faster computational times. They were particularly impressed with Seokkwan Yoon's presentation of the performance of the INS3D-LU code. By April 1990, work with INS-3D was suspended and I was using the LU code full time.

To the LU code I added periodic boundary conditions, grid singularity conditions, and a turbulence model. The periodic conditions required conversion of several subroutines to periodic form. I also optimized the code to run as fast as 6 μ seconds/point/iteration on one Cray Y-MP processor.

In order to extract data required for the SUBOFF project, I wrote a general interpolation code, which returns the Q variables for any x,y,z location within a grid, using a tetrahedral decomposition of the grid. The results were presented at a SUBOFF meeting in Annapolis⁹.

The interest of the SUBOFF committee was limited to a small portion of the requested data, which happened to include the regions most poorly modeled by our grids and the LU code. As a result, the work fared poorly in SUBOFF evaluations.

On return to Ames, I set out to find the true source of the errors, using test cases which would isolate potential problems. Periodic flows about a cylinder and flow about a hemisphere cylinder with a grid singularity converged nicely. A coarse grid on a flat plate reproduced the Blasius boundary layer for low Reynolds numbers. The only case with exceptional errors was one which attempted to compute a high Reynolds number solution on a flat plate. After extensive tests of various grid stretching strategies and smoothing parameters, this problem remained. The deficiency of the LU code in convergence on fine grids has still not been corrected, although significant progress has been made recently by Goetz Klopfer and Dean Kontinos.

In order to continue with the SUBOFF project, coarsened grids were used. With spacing of 1.0E-3, the solution converged, but with no boundary layer development at all. With 1.0E-5 spacing, an approximation to the boundary layer was obtained, but convergence was poor. Aside from inaccurate boundary layer thickness, a reasonable representation of the flow physics was obtained.

The improved results were presented to SUBOFF committee members at David Taylor Research Center in September 1990. They were pleased by visualizations of vortex separation phenomena near the fairwater which other SUBOFF participants had missed. As expected, wake-survey data showed only qualitative agreement with experimental results, due to the inability to converge solutions on a fine grid.

The SUBOFF committee appeared to conclude that no available CFD code was adequate for outright prediction of drag or moments on a full submarine. Funding originally planned for ongoing work was redirected, and the project ended. A promised final report was never distributed.

Other work during this time period included testing of non-reflective boundary conditions based on the work of Bayliss, Gunzburger, and Turkel. Although developed for spherical outer boundaries and laminar flow, the conditions improved convergence on the SUBOFF cases, and allowed an outer boundary only two body lengths from the vehicle surface. I served as mentor for highschool student Adam Nash, and he was helping me make test runs to quantify the effect of the nonreflective boundary conditions when this

project ended in December 1990.

Parallel CFD Computation

The national HPCCP program was established, including a Computational Aerosciences (CAS) element. CFD was to be performed on highly parallel computers, with faster testbed machines obtained on a regular basis. Long-range goals included teraflops-rate multidisciplinary optimization of aerospace vehicles.

In January 1991, I began exploring the use of parallel, distributed memory computers for CFD. The initial plan was to port the CNS (Compressible Navier-Stokes) code to the Intel iPSC/860. CNS was a code which I developed for my Ph.D. thesis work in 1987 and 1988^{1,2}, and which became a standard within the Applied Computational Fluids (later Computational Aerosciences) branch at NASA Ames.

My first tests on the iPSC/860 involved a simple thermal relaxation code, which was small enough to allow short compilation times and easy modification. This code later became instrumental in testing parallel I/O strategies, and was requested by researchers at NAS, RIACS, Dartmouth, and several Intel Corporation sites. This work revealed some surprising behavior of the iPSC/860 I/O system, and was used in defining I/O requirements for future machines¹⁰.

The CNS port was under way from March through August 1991, when I learned that Sisira Weeratunga had parallelized the ARC3D algorithm on the iPSC/860. His work was based on the OVERFLOW code, and he was using a more sophisticated parallelization method, so I dropped the CNS port and got a copy of what later was called Parallel OVERFLOW. I added parallel I/O to the code (it could only run a single, internally-generated grid), as well as a restart capability. This allowed testing to begin with a flat plate case, which ran at 32 microseconds/point/iteration on 32 processors.

In October 1991, I computed an Euler solution on the Boeing 1807 wing body, a preliminary HSCT (High Speed Civil Transport) design. For this case, computational times improved to 28.9 μ sec/pt/it on 32 processors, with timings of 18.8 and 11.1 on 64 and 128 processors, respectively. Several tests were made to explore differences between the parallel solution and UPS Cray results. The differences were identified to be the result of different effective grid resolution in the two codes.

In November 1991, Ron Bailey identified my wing-body results (presented by Tom Edwards) as possibly the first 3-D external flow calculations on a massively parallel machine. This led to a video¹¹ describing the work, later used in HPCCP program reviews by Ken Stevens, and finally incorporated in a professionally-developed video used for presentations to Congress during the budgeting process. Graphics from this project have been requested dozens of times, and included in the publication "High Performance Computing and Communications: Toward a National Information Infrastructure", otherwise known as the "FY 1994 Blue Book" of the National Coordination Office for HPCC.

In March 1992, I completed a parallel Baldwin-Lomax model. This allowed Navier-Stokes HSCT solutions to be computed in parallel¹³. An attempt to present the work at "Parallel CFD '93" was thwarted by the 427 approval process, which took 120 days, extending past the abstract deadline. I refused an invitation to present the work at "Physics Computing '93" in Albuquerque, New Mexico, on the recommendation of my technical monitor, due to joint sponsorship by the European Physical Society.

In July 1992, work began with a new version of Parallel OVERFLOW, in which Sisira Weeratunga had included Chimera grid capabilities. This allowed the computation of more complex geometries, limited only by grid generation capabilities and available memory on the parallel computers. Computational time was now down to about 20 μ seconds/point/iteration on 32 processors. Computations included a wing/nacelle in September 1992, wing/body/nacelle in October 1992, and wing/nacelle/diverter in October 1993^{4-6,14,15}.

In response to a Boeing request in April 1993, I prepared a test case for the Parallel OVERFLOW code, and assisted John Wai of Boeing Military Aircraft in running the case on their parallel system.

Throughout this project I have contributed to other's work by enhancing communications. I maintained an email list and sent written notes of NAS Parallel Systems meetings to interested RFA civil service and contractor researchers. At those meetings I was often the only user representing our needs to the NAS staff. I established a series of meetings on HPC and CAS topics. Half of the meetings were for MCAT Parallel Computing Section employees only, and are continuing. The others were open technical sessions, which were converted to branch HPC meetings by Terry Holst.

I contributed to the K-12 education mandate of HPCCP, by mentoring a teacher during the summer of 1993. This led to a package of teaching materials, in both paper and electronic form, which introduces the concept of representing continuous quantities on a discreet grid. I also organized a 3-hour session to introduce a group of high school teachers from Mendocino County to concepts which might be introduced in their classrooms.

Parallel Aerodynamic Optimization

In May 1993, I first ran Parallel OVERFLOW under the control of the NPSOL optimizer, moving closer to the design goals of HPCC. This tool was tested against the Haack-Adams results of Samson Cheung and Phil Aaronson, and later applied to a wing-fuselage geometry. The work required new parallel grid-modification routines, and parallelization of force and moment routines. The code also had to be modified to accomodate larger amounts of Chimera interface information than had been provided for.

The work using NPSOL used parallel computation within each flow solution, but ran those solutions one at a time. This meant that an additional level of parallelism was going unexploited. NPSOL was difficult to parallelize, and in February 1994, Samson Cheung wrote a simpler version of the quasi-Newton optimization method which was designed for parallelization. I developed a grid generation code which would produce wing surface and volume grids as a function of a broad set of design variables. These include span, chord, twist, sweep, thickness, and camber. Each variable can be described multiple locations where appropriate.

The combined wing generation and optimization code was originally targeted for the iPSC/860 or the Intel Paragon, but announcement of the impending acquisition of an IBM SP-2 system changed the focus to workstation clusters which have more in common with the SP-2. The code was tested on the NAS SPS machines, and on the RFA workstation cluster, with encouraging results. The required cases for derivative calculations and line searches were done in parallel, using one serial flow solver on each workstation. Maximum parallel speedup will be achieved by using the MEDUSA code of Merritt Smith for distribution of zones across multiple machines, and the multipartitioning algorithms of Rob van der

Wijngaart for parallelism within zones.

Work Since August 1994

This time period began with efforts to extend prior results to new architectures. The OVERFLOW solver was compiled on the IBM SP2 computer, and an account on a Cray T3D at Pittsburg was obtained. On the SP2, I worked with the PVM and MPI message passing libraries.

A major goal for this period was to compute a 19-zone, 5 million point Navier-Stokes solution about the Boeing Reference H geometry with engine nacelles and diverters. Such a large problem would be impossible on the iPSC/860, so an account on the Intel Paragon was obtained. A copy of Weeratunga's parallel flow solver proved to be poorly debugged, relative to prior versions on the iPSC/860. I identified and fixed 7 bugs which prevented the code from being used effectively. I was then able to obtain two to three orders of magnitude convergence on a 7-zone subset of the target problem. Poor Chimera interfaces are the likely reason that further convergence is difficult. The nacelles were not added during the cooperative agreement period.

Graphics of my optimization and HSCT work were in demand during this period. I provided slides to Terry Holst, Merritt Smith, Guru Guruswamy, and I.C. Chang for a variety of presentations. I also supported a method for extracting cross-sectional data from CFD datasets, which I developed during HSCT validation studies. Customers included Francisco Torres and Joseph Garcia. Finally, I did some educational outreach by visiting a classroom of 4th graders preparing to attend the Ames Aerospace Encounter, and by answering questions for a high school student who has a grant to write a paper on NASA policy.

Summary and Conclusions

During the last five years, CFD has matured substantially. Pure CFD research remains to be done, but much of the focus has shifted to integration of CFD into the design process. The work under these cooperative agreements reflects this trend. The recent work, and work which is planned, is designed to enhance the competitiveness of the US aerospace industry. CFD and optimization approaches are being developed and tested, so that the industry can better choose which methods to adopt in their design processes. The range of computer architectures has been dramatically broadened, as the assumption that only huge vector supercomputers could be useful has faded. Today, researchers and industry can trade off time, cost, and availability, choosing vector supercomputers, scalable parallel architectures, networked workstations, or heterogenous combinations of these to complete required computations efficiently.

References

The lists below include minor publications and talks as well as those presented at major conferences or in journals. The work often involved sensitive technology, so that presentation in closed meetings was preferred.

Publications

¹James S. Ryan, "Application of Computational Fluid Dynamics to Three-Dimensional Bodies in Hypersonic Flow," Ph.D. Thesis, University of Colorado, Boulder, December 1989.

²James S. Ryan, Jolen Flores, and Chuen-Yen Chow, "Development and Validation of a Navier-Stokes Code for Hypersonic External Flow," *Journal of Spacecraft and Rockets*, Vol. 27, No. 2, 1990, pp. 160-166.

³Seokkwan Yoon, Stuart Rogers, James S. Ryan, Cetin Kiris, Moshe Rosenfeld, and Dochan Kwak, "Numerical Methods for Three-Dimensional Incompressible Navier-Stokes Equations," NASA/OAST CFD Conference, NASA Ames Research Center, March 1990.

⁴James S. Ryan, Sisira Weeratunga, and Samson Cheung, "Parallel Navier-Stokes Computation of Supersonic Vehicle Flowfields," Compendium of Abstracts, NASA Computational Aerosciences Conference, August 18-20, 1992.

⁵James S. Ryan and Sisira Weeratunga, "Parallel Computation of 3-D Navier-Stokes Flowfields for Supersonic Vehicles," AIAA Paper 93-0064, Reno, NV, January 1993.

⁶Philip G. Aaronson and James S. Ryan, "OVERFLOW Code Validation Study for the Boeing Reference H," Proceedings of the First NASA/Industry High Speed Research Propulsion/Airframe Integration Workshop, Cleveland, OH, October 26-27, 1993.

Presentations

⁷James S. Ryan "Submarine Flow Simulation," SUBOFF meeting at NASA Ames, December 12, 1989.

⁸James S. Ryan "Naval Hydrodynamics," RFA Branch Review at NASA Ames, February 26, 1990.

⁹James S. Ryan "Submarine Hydrodynamics," SUBOFF Review Meeting, United States Naval Academy, Annapolis, MD July 11-12, 1990.

¹⁰James S. Ryan, "Experience with Fluid Dynamics Data on the iPSC/860 Concurrent File System," presented at the Concurrent I/O SIG at Supercomputing '91, Albuquerque, NM, November 1991.

¹¹James S. Ryan, "A Fluid Dynamics Module for HPCCP Multidisciplinary Optimizations," a video, December 5, 1991. Portions were used in "Grand Challenges 1993," produced by James H. Cosgrove for NASA Headquarters.

¹²James S. Ryan, "High Speed Civil Transport Aerodynamics," presented at a review for Lee Holcomb of NASA Headquarters, January 29, 1992.

¹³James S. Ryan, "Progress on HSCT Applications," presented at the NASA Computational Aerosciences Industry Briefing, Cleveland, OH, June 18, 1992.

¹⁴James S. Ryan, "Parallel Computation of Supersonic Flowfields," presented at the Code R CFD Peer Review at NASA Ames, November 19, 1992.

¹⁵James S. Ryan, "HSCT Grand Challenge," RFA Branch Review at NASA Ames, March 24, 1993.